arXiv:1406.2852v2 [cs.DC] 16 Jun 2014

On the Impact of Geometry on Ad Hoc Communication in Wireless Networks *

Tomasz Jurdzinski[1] Dariusz R. Kowalski[2] Michal Rozanski[1] Grzegorz Stachowiak[1]

Abstract

In this work we address the question how important is the knowledge of geometric location and network density to the efficiency of (distributed) wireless communication in ad hoc networks. We study fundamental communication task of broadcast and develop well-scalable, randomized algorithms that do not rely on GPS information, and which efficiency formulas do not depend on how dense the geometric network is. We consider two settings: with and without spontaneous wake-up of nodes. In the former setting, in which all nodes start the protocol at the same time, our algorithm accomplishes broadcast in $O(D \log n + \log^2 n)$ rounds under the SINR model, with high probability (whp), where D is the diameter of the communication graph and n is the number of stations. In the latter setting, in which only the source node containing the original message is active in the beginning, we develop a slightly slower algorithm working in $O(D \log^2 n)$ rounds whp. Both algorithms are based on a novel distributed coloring method, which is of independent interest and potential applicability to other communication tasks under the SINR wireless model.

1 Introduction

In this paper we study distributed communication problems in wireless networks, where interferences are resolved by the Signal-to-Interference-and-Noise Ratio (SINR) physical model. Specifically, but not exclusively, we concentrate on the broadcast problem, where a piece of information stored in a specified station/node (the source) is supposed to be delivered to all other stations in the network. The broadcast is a fundamental communication primitive, whose complexity is well understood in the previous models of wireless communication, such as radio networks. Closer to reality, models based on the SINR constraint attracted attention of algorithmic community much later than the radio network model. One of the key differences between models is that radio networks take into account only interference between stations in close neighborhood, while the SINR model relies on the physical assumptions that the strength of signals decrease gradually according to a continuous function and cumulate, which makes development of algorithms and their analysis much more complicated.

In this work we show that efficiency of wireless communication depends mainly on parameters of the communication graph, even for devices with limited knowledge and capabilities. In particular, we do not assume any carrier sensing capabilities, initial synchronization or any knowledge other than rough estimates of the number of nodes and physical SINR parameters. Despite of that, we develop almost optimal and well scalable solutions to the broadcast and wake-up problems. Moreover, as mentioned above, the worst-case performance of the considered communication tasks, and other problems building on them, depends only on the topology of the communication graph, also called a reachability graph, and not on specific location of nodes within reachability balls. One of the implications is that geometric properties of reachability regions studied in some previous works, c.f., [1], do not influence worst-case scenarios in ad hoc communication by more than $O(\log^2 n)$ factor - the factor by which our algorithms are far from lower bounds.

[1]Institute of Computer Science, University of Wrocław, Poland.
[2]Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK.
*This work was supported by the Polish National Science Centre grant DEC-2012/07/B/ST6/01534.

1.1 Model

We consider the model of a wireless network consisting of *stations*, also called *nodes*, deployed into a metric space with *bounded growth property* of degree γ.[1] All stations are identical, and therefore each of them has the same transmission power P (we call it a *uniform power model*).

There are three fixed model parameters: path loss $\alpha > \gamma$, threshold $\beta \geq 1$, ambient noise $\mathcal{N} > 0$. We also assume a connectivity graph parameter $\varepsilon \in (0, 1)$.

The $SINR(v, u, \mathcal{T})$ ratio, for given stations u, v and a set of (transmitting) stations $\mathcal{T}$, is defined as follows:

$$SINR(v, u, \mathcal{T}) = \frac{P\text{dist}(v, u)^{-\alpha}}{\mathcal{N} + \sum_{w \in \mathcal{T} \setminus \{v\}} P_w\text{dist}(w, u)^{-\alpha}} \tag{1}$$

In the *Signal-to-Interference-and-Noise-Ratio (SINR) model* a station u successfully receives a message from a station v in a round if $v \in \mathcal{T}$, $u \notin \mathcal{T}$ and $SINR(v, u, \mathcal{T}) \geq \beta$, where $\mathcal{T}$ is the set of stations transmitting in that round.

Synchronization It is assumed that algorithms work synchronously in rounds. In general, we do not assume global clock ticking. Note, however, that some kind of global synchronization can be achieved by appending a counter to every message sent throughout broadcast algorithm; we use this property in algorithms developed in this work.

Carrier sensing We consider the model *without carrier sensing*, that is, a station u has no other feedback from the wireless channel than receiving or not receiving a message in a round t.

Knowledge of stations Each station knows the number of stations in the network, n. Our algorithms also work when stations share, instead of n, an estimate $\nu \geq n$ of this value which is $O(n^c)$ for a fixed constant c. We assume that nodes do not know the precise value of the SINR parameters α, β, and $\mathcal{N}$ but instead know only upper and lower bounds for the parameters (i.e., $\alpha_{\min}$ and $\alpha_{\max}$, $\beta_{\min}$ and $\beta_{\max}$, $\mathcal{N}_{\min}$ and $\mathcal{N}_{\max}$). For simplicity, in this version of the paper we perform calculations assuming that exact values of these parameters are known. In order to take into account uncertainty regarding those parameters, it is sufficient to choose their maximal/minimal values depending on the fact whether upper or lower estimates are provided.

Messages and initialization of stations We consider two variants of initialization of stations: without spontaneous wake-up and with spontaneous wake-up. In the former model each station (except some distinguished one(s)) sleeps till it obtains the message for the first time. In the latter variant, all nodes are woken up at the same time and start an execution of an algorithm simultaneously. (Observe that nodes can benefit from the spontaneous wake-up setting by performing a local preprocessing simultaneously for the whole network.)

Each station can either act as a sender or as a receiver during a round. A sender can transmit a broadcast message with attaching to it $O(\log n)$ additional bits. Our algorithms are described from a "global" perspective, i.e., we count rounds starting from the moment when the first message is sent. In order to synchronize stations in the model with non-spontaneous wake-up, we assume that each message contains the number of rounds elapsed from the beginning of the execution of the algorithm.

Ranges and uniformity The *communication range* r is the radius of the ball in which a message transmitted by a station is heard, provided no other station transmits at the same time. Note that $r = (P/(\mathcal{N}\beta))^{1/\alpha}$, where P is the transmission power of a station, c.f., Equation (1). Without loss of generality we assume that $r = 1$. (Note that this assumption implies the relationship $P = \mathcal{N}\beta$.)

Communication graph and graph notation The *communication graph* $G(V, E)$ of a given network consists of all network nodes and edges (v, u) such that $\text{dist}(v, u) \leq (1 - \varepsilon)r = 1 - \varepsilon$, where $0 < \varepsilon < 1$ is a fixed model parameter. The meaning of the communication graph is as follows: even though the idealistic communication range is r, it may be reached only in a very unrealistic case of single transmission in the

[1] This notion generalizes the Euclidean R^γ space; its formal definition is provided later in this section.

whole network, c.f., [5]. In practice, however, many nodes located in different parts of the network often transmit simultaneously, and therefore it is reasonable to assume that we may only hope for a slightly smaller range to be achieved. The communication graph, through restricting connections to ranges at most $1-\varepsilon$, envisions the network of such "reasonable reachability". It has become a classic tool in the analysis of ad hoc communication tasks under the SINR physical model c.f., [5, 14, 21].

Note that the communication graph is symmetric for uniform networks. By a *neighborhood* of a node u we mean the set of all neighbors of u in G, i.e., the set $\{w \mid (w,u) \in E(G)\}$. The *graph distance* from v to w is equal to the length of a shortest path from v to w in the communication graph, where the length of a path is equal to the number of its edges. The diameter D of a network is equal to the diameter of its communication graph (i.e., the largest graph distance between any pair of nodes), provided the graph is connected.

Metric space Given a metric space with a distance function dist, $B(v,r)$ for a point v from the space and $r > 0$ is equal to $\{w \mid \text{dist}(v,w) \leq r\}$ and is called a *ball* with radius r and center v. A *unit ball* is a ball with radius 1. Moreover, let $\chi(a,b)$ denote the number of balls with radius b sufficient to cover a ball with radius a. Nodes of a network are embedded (as points) in a general metric space with a distance function dist that satisfies the following *bounded growth property*: For every $d > 0$, $c \in \mathbb{N}$ and a point v in the metric space, the ball $B(v, c \cdot d)$ is included in a union of $O(c^\gamma)$ balls with radius d, where γ is a parameter called a *dimension* of the metric. (That is, $\chi(cd, d) = O(c^\gamma)$ for each $d > 0$ and $c \in \mathbb{N}$.) Note that this in particular implies that $B(v,(c+1)\cdot d) \setminus B(v, c \cdot d)$ can be covered by $O(c^{\gamma-1})$ balls with radius d; we will often rely on this property in our analysis when estimating the total strength of the interference received at a node.

Broadcast problem In the broadcast problem, there is one distinguished node, called the *source*, which initially holds a piece of information (also called a *source message* or a *broadcast message*). The goal is to disseminate this message to all other nodes in a network with connected communication graph. We are interested in minimizing the *time complexity* of this task being the minimum number of rounds after which, for all communication networks defined by some set of parameters, the broadcast occurs with high probability. This time is counted since the source is activated.

1.2 Previous and related work

The algorithmic research on communication in the SINR networks started around 10 years ago. Most papers concentrate on one-hop communication, which includes the local broadcast problem [9, 11, 22], link scheduling [18, 10], connectivity [2, 12] and others. Among them, the most related to this work are papers on local broadcast, in which each node has to transmit a message only to its neighbors in the corresponding communication graph. Using the local broadcast algorithm (e.g., from [11]) as a building block yields a solution for (global) broadcast that runs in $O(D(\Delta + \log n)\log n)$ time, where Δ is the maximal degree of the communication graph. However, since there is only one message to be propagated in the global broadcast, we would like to avoid the dependence on potentially large parameter Δ (which could be necessary when all senders have different messages, but not in the case of a single source global broadcast).

In order to address obstacles for multi-hop communication, various authors take advantage of several features helping to design efficient algorithms. As for the broadcast problem in the SINR model, Scheideler at al. [20] solve the problem in $O(D + \log^2 n)$ rounds using a tunable collision detection and assuming that all stations start a protocol simultaneously, which allow them to build an overlay structure along which the message is then propagated. Yu et al. [21] solve the problem in $O(D + \log^2 n)$ rounds using power control, allowing stations to decide the strength of a transmitted signal in each step. Moreover, their results works merely for a restricted family of networks, excluding the most challenging scenarios. Specifically, their algorithm works under assumption that, for each node v, its closest neighbor is in distance at most $1/3$. Moreover, a possibility of filtering out messages received from large distances is necessary.

In [14] an $O(D\log n + \log^2 n)$ randomized algorithm and in [15] an $O(D\log^2 n)$ deterministic algorithm for networks deployed in the Euclidean space are presented, where stations know their own positions (e.g., thanks to GPS devices). Finally, Daum et al. [5] designed an algorithm working in $O((D\log n)\log^{\alpha+1} R_s)$ rounds, provided stations know only granularity R_s of the network (i.e., the maximum ratio between actual

distances of stations connected by an edge in the communication graph) and do not use any other additional features.

If the model prevents successful direct transmissions between nodes which are not connected in the communication graph, so called *weak device model*, the lower bound $\Omega(D\Delta)$ holds even if stations know their positions on the plane [16], which separates that model from the one considered in this work. For other related problems in this more harsh model see e.g., [9, 13].

In a related radio network model, the complexity of broadcasting is much better understood. Its complexity in the model without collision detection is $\Theta((D + \log n)\log(n/D))$ [4, 19]. Interestingly, this lower bound was recently broken for the model with collision detection [8], in which a solution in $O(D + \text{polylog}(n))$ was designed. For the easier case where all nodes start during the same round, it is currently unknown whether or not formulas better than the ones in general graphs could be obtained, but in unit disk graphs a solution of the form $O(D + \log^2 n)$ is likely possible [5]. As shown e.g. in [7, 6], geometric graphs exhibit more efficient solutions than those possible in a general graph model of radio networks.

1.3 Our results

The results of this paper state that the broadcast problem can be accomplished in $O(D\log^2 n)$ rounds in the model without spontaneous wake-up and in $O(D\log n + \log^2 n)$ rounds in the model with spontaneous wake-up. Interestingly, this performance formulas does not depend on any geometric parameter related with specific locations of nodes, only on parameters of communication graph (expressing the relation whether nodes are within their transmission ranges or not). This improves the result of Daum et al. [5] for $R_s = \omega(2^{(\log n)^{1/(1+\alpha)}})$, where R_s is the maximum ratio between between actual distances in the metric space of stations which are connected by an edge in the communication graph. (Thus, in particular, for $R_s = \Omega(n^\delta)$, for any fixed $\delta > 0$.) Moreover, our algorithm does not need information about parameter R_s. As R_s might be even exponential wrt to n,[2] this is the first solution with guaranteed $O(D\text{polylog}(n))$ complexity in the SINR networks without spontaneous wakeup, power assignment, carrier sensing (tuned collision detection), or any knowledge about location of nodes.

As the main tool, we design a specific coloring algorithm associating with each active station v the probability p_v which, when fixed, helps to solve other communication problems efficiently, including the consensus problem, the leader election and the alert protocol problem. This coloring plays a role similar to backbone structures in many other communication models (c.f., [13, 23]). In order to get rid of the dependence on the granularity parameter R_s, we use a different approach to [5]. They allow all stations to transmit with constant probability, which may generate a lot of noise but makes possible communication between stations within the smallest distance in the network. We, on the other hand, start from very low probabilities of the order $1/n$ and increase them gradually until stations can hear reasonable number of messages. Such a strategy reminds solutions to the local broadcast problem (e.g., in [11]); however, unlike in those solutions (which did not need it), the key and subtle issue in our approach is to somehow distinguish by a station v between the densities of the network in close neighborhood $B(v, \varepsilon/2)$ and in broader neighborhood $B(v, 1)$. In order to tackle this issue without any geolocation information and other tools such as power control, we proceed by interleaving two kinds of phases serving different purposes: phases where stations transmit with some assigned probabilities (which intuitively grow up gradually) with phases where probabilities of transmissions are "scaled up" based on local statistics of successful transmissions with carefully probed transmission probabilities. This approach faces various technical obstacles, mainly due to the lack of geolocation information, which are addressed in the paper.

Organization of the paper Basic properties of simple transmission scenarios are given in Section 2. The main coloring tool, its details and construction can be found in Section 3. Its applications to broadcasting in non-spontaneous and spontaneous settings are presented in Section 4. Missing proofs can be found in the full version of the paper.

[2] Consider, for example, n stations $x_1, \ldots, x_n$ on a line such that $\text{dist}(x_i, x_{i+1}) = 1/2^i$.

2 Notations and Technical Preliminaries

We say that an event happens in a network of n stations *with high probability (whp)* when the probability is at least $1 - 1/n^c$, for some constant $c > 0$.[3] An event occurs with *negligible* probability if its negation occurs whp. In particular we prove in this paper that our algorithms succeed whp.

Given a metric space with the distance $\text{dist}(\cdot, \cdot)$, we use the notation $B(v, r) = \{v' \,|\, \text{dist}(v, v') \leq r\}$. For a set of stations A, by $N(A)$ we denote the set of their neighbors in the communication graph, i.e., $N(A) = \{w \,|\, \text{dist}(v, w) \leq 1 - \varepsilon \text{ for some } v \in A\}$.

In order to simplify calculations, we assume that the constant hidden in the expressions $O(c^\gamma)$ determining the growth parameter of the metric space is equal to 1 (this does not change the asymptotic complexity of our algorithms).

Below, we formulate a basic property that a station transmitting successfully to a distance larger than $1-\varepsilon$ delivers its message to the neighbors of other stations in its close proximity.

Fact 1. *If a station v is transmitting in a round and its message can be successfully received at each point u such that dist$(u, v) \leq 1 - \varepsilon/2$, then the message is received by all neighbors of all nodes from $B(v, \varepsilon/2)$.*

Given a set of stations T transmitting in a round and a station u, the *interference* at u is equal to $I_u = \sum_{w \in T \setminus \{v\}} P/(\text{dist}(u, w))^\alpha$, where v is a station in the smallest distance from u among the elements of T.

Fact 2. *Let $x \leq 1/2^{1/\alpha}$. If the interference at some receiver u is at most $\mathcal{N}/(2x^\alpha)$, then it can hear the transmitter v from the distance x.*

Proof. Let us recall assumption that $P = \mathcal{N}\beta$. We have

$$SINR(v, u, T) \geq \frac{\beta\mathcal{N}}{x^\alpha}/(\mathcal{N} + \frac{\mathcal{N}}{2x^\alpha}) = \beta/(x^\alpha + 1/2) \geq \beta \,,$$

where the former inequality follows from the bound on the interference and the latter from the assumption $x \leq 1/2^{1/\alpha}$. □

Fact 3. *If the interference at some receiver u is at most $\mathcal{N}\alpha x$, then it can hear the transmitter v from the distance $1 - x$.*

Proof. By the Bernoulli inequality we get $(1 + x)^\alpha \geq 1 + \alpha x$. Thus

$$SINR(v, u, T) \geq \frac{P/(1-x)^\alpha}{\mathcal{N} + \mathcal{N}\alpha x} \geq \frac{P}{\mathcal{N}(1+x)^\alpha(1-x)^\alpha} = \frac{P}{\mathcal{N}(1-x^2)^\alpha} \geq \frac{P}{\mathcal{N}} = \beta \,.$$

□

Consider a scenario where every station v is assigned a variable p_v being its transmission probability.

Fact 4. *[17] Assume that $\sum_{v \in A} p_v = s \leq 1/2$ for some set of stations A. Then the probability that exactly one element of A transmits is at least $s/2$ and at most s.*

Fact 5. *[17] Let $p_v \leq 1/2$ for every station $v \in A$. Then the probability that no station from A transmits is at least $(1/4)^{\sum_{v \in A} p_v}$.*

We say that *bounded density property* is satisfied with the parameter $C > 0$ if $\sum_{w \in B} p_w \leq C$ for every unit ball B. The *effective communication property* is satisfied if the probability that a station v hears w when w is the only transmitting station in $B(v, 2/3)$ is at least $1/2$.

Fact 6. *For any network parameters $\alpha > \gamma$, $\beta \geq 1$, $\mathcal{N} > 0$ and $\varepsilon < 1$, there exists a constant $C > 0$ such that if the bounded density property is satisfied with any parameter $C_1 < C$, then the effective communication property is satisfied as well.*

[3] We often use union bounds to show that some undesirable events happen with small probability. Therefore, in order to carry on we usually require for some basic events occurring during the analysis to happen with probability at least $1 - 1/n^4$, when saying that they occur whp. This requirement will guarantee that after applying all union bound arguments within the analysis, the probability that our algorithms accomplish their tasks correctly within the given time bounds is at least $1 - 1/n$.

Proof. Assume that w is the only transmitter in $B(v, 2/3)$ and $\sum_{w \in B} p_w \leq C$ for every unit ball B and a constant C. Then, v receives the message from w by Fact 3, provided the interference from the remaining area is smaller than $\mathcal{N}\alpha \cdot \frac{1}{3}$. The expected value of this interference under the bounded density property with the parameter C is

$$E(I_v) \quad \leq \quad C \sum_{i>0} \frac{P \cdot O(i^{\gamma-1})}{\mathcal{N} \cdot (\frac{2}{3} i)^{\alpha}} = \frac{CP(3/2)^{\alpha}}{\mathcal{N}} O(\sum_{i>0} i^{\gamma-\alpha-1}) \leq C \cdot C'$$

where C' is a constant depending on γ, α, β, $\mathcal{N}$. The first among the above inequalities follows from the bounded growth property and the last one from the assumption that $\gamma < \alpha$. Thus if $C < \frac{\frac{1}{6}\mathcal{N}\alpha}{C'}$ then $E(I_v) \leq C \cdot C' < \frac{1}{2} \cdot \frac{1}{3}\alpha$. Using Markov bound, we get $P(I_v > \mathcal{N}\alpha \cdot \frac{1}{3}) < 1/2$ and the probability that v receives a message from w is at least $1/2$ by Fact 3. □

3 Network Coloring

The key ingredient of broadcasting algorithms presented in this paper is the procedure StabilizeProbability (Algorithm 1), which assigns probability ("color") p_v from the set

$$\{2^i p_{\text{start}} \mid i \in [0, \lceil \log(p_{\max}/p_{\text{start}}) \rceil\}$$

to each station v participating in an execution, where $p_{\text{start}} = \Theta(1/n)$ and $p_{\max}$ is a constant which will be specified later. Thus, the number of colors is $O(\log n)$.

Before giving details of the procedure StabilizeProbability, we state the key properties that we want the procedure to satisfy. We express them as properties of the obtained coloring, described by the following lemmas. These lemmas are true for some constants C_1 and C_2 that depend on ε, γ, the parameters of the SINR model and constants chosen in the algorithm.

Lemma 1. *After an execution of StabilizeProbability for a set of stations A, the inequality*

$$\sum_{\substack{w: p_w = p \\ w \in B}} p_w < C_1$$

holds for every color p and unit ball B whp.

Lemma 2. *After an execution of StabilizeProbability on a set of stations A, for every $v \in A$ there exists a color p such that the following inequality holds whp:*

$$\sum_{\substack{w: p_w = p \\ w \in B(v, \varepsilon/2)}} p_w \geq C_2 \, .$$

In this section we describe formally the algorithm StabilizeProbability and prove that it satisfies Lemma 1 and Lemma 2 for appropriate constants C_1 and C_2. Its pseudo-code is given as Algorithm 1. The pseudocode is missing information about the actual values of constants c_0, c_1, c_2, c_3, c', c_ε, C_2, and p_{start} used in the algorithm. We will choose the appropriate values for those constant in the analysis of the algorithm and its properties.

Algorithm StabilizeProbability performs two kinds of tests, defined by sub-routines DensityTest and Play-off, each of them taking $O(\log n)$ rounds. The while loop of StabilizeProbability is repeated $O(\log n)$ times, since each node v starts with $p_v = \Theta(1/n)$, increases p_v twice in each repetition of the loop and finishes either in line 6 or after achieving $p_v \geq p_{\max} = \Theta(1)$. Therefore, the following claim holds.

Fact 7. *Algorithm StabilizeProbability works in $O(\log^2 n)$ rounds whp.*

In Section 3.1, we give a general idea why StabilizeProbability satisfies Lemma 1 and Lemma 2.

In Section 3.2, the constants c_0, c_1 in the sub-routine DensityTest are set so that this procedure helps each station v to estimate whether density (i.e., the sum of probabilities assigned to stations) in $B(v, 1)$ already achieved constant value.

Proofs of Lemma 1 and Lemma 2 are presented in separate sections. In Section 3.3, containing the proof of Lemma 1, the constants c_2, c_3 for sub-routine Playoff are chosen in order to guarantee that the regions (unit balls) with largest density will become sparser, i.e., some stations switch off (line 6 of Algorithm 1) in them in each execution of Playoff whp. Combined with the appropriate choice of the constant c' from Playoff, this will assure that Lemma 1 is satisfied.

In Section 3.4, dedicated to the proof of Lemma 2, the constant c_ε is chosen in order to make very unlikely the situations in which a station v switches off when the sum of probabilities of active stations in $B(v, \varepsilon/2)$ is very small. This property, combined with the fact that probabilities of "active" stations (i.e., stations which are not switched off in line 6 of Algorithm 1) grow up to the constant $p_{\max}$, will lead to the statement of Lemma 2.

3.1 Overview of the algorithm

The coloring algorithm assigns probability/color p_v to each active station $v \in A$ such that there are at most $\log n$ various colors and simultaneously the following two properties hold:

1. For each color, the sum of probabilities of stations in this color in each unit ball is at most C_1, for some constant C_1 (Lemma 1).

2. For each active station v, there exists a color such that the sum of probabilities of this color in the ball $B(v, \varepsilon/2)$ is at least C_2, for some constant C_2 (Lemma 2).

The former property assures that, when all stations transmit with assigned probabilities, the expected interference coming from the whole network is small at any station (this follows from the assumption that $\alpha > \gamma$, which implies $\sum_{i \in \mathbb{N}} i^{\gamma-1-\alpha} = O(1)$). Thanks to this property, if a station v is the only transmitter in $B(v, 2d)$ for a constant $d \leq 1 - \varepsilon$ then the message transmitted by v can be received in each point of $B(v, d)$ with constant probability. The latter property on the other hand guarantees that for each station v, the probability that a station from $B(v, \varepsilon/2)$ transmits is constant as well, for some color. Both properties combined imply that, each station from $N(A)$ receives a message with probability $\Omega(1/\log n)$ in a round if each station $v \in A$ transmits with probability $p_v/\log n$.

There is some intuition behind the coloring algorithm. The optimal probability for a station v to transmit is approximately $1/|\{w : \text{dist}(v, w) \leq \varepsilon/2\}|$. This would ensure that the sum $\sum_{B(v,\varepsilon/2)} p_v$ is limited from below by constant C_2 and C_1 is of the order of C_2/ε^γ. The algorithm starts from low probabilities (smaller than C_1/n), continuously increasing them. Once a station v starts receiving messages from others, it assumes that the sum of probabilities in $B(v, 1)$ is constant, which indicates that further increase of all probabilities in $B(v, 1)$ could break the property 1 (Lemma 1). When applied to the task of local broadcasting [11], this means that the "right" probability is reached (so, probabilities could be frozen). If stations are uniformly distributed, then this implies that the average sum of probabilities in balls of diameter ε is around C_1/ε^γ, which would also satisfy the property 2 (Lemma 2) for $C_2 \approx C_1/\varepsilon^\gamma$. In general, the constant sum of probabilities, around C_1, of stations in $B(v, 1)$ does not exclude that the sum of probabilities in $B(v, \varepsilon/2)$ is still very small.[4] Hence, our intuitive goal at this stage would be to distinguish those regions (balls of diameter ε) in which the sum of probabilities exceeds (say) half of the average from those where it is much smaller.

The main difficulty is to sense the actual sum of probabilities in $B(v, \varepsilon/2)$, without possibility of filtering out messages received from distance larger than $\varepsilon/2$ (as there is no geolocation). Imagine that stations in a network are uniformly distributed and the probabilities have reached such values that the sums in a unit ball are close to C_1. If one replaces p_v with $c_\varepsilon p_v$ for large enough c_ε (depending on the growth parameter of the metric, e.g., $c_\varepsilon \approx 1/\varepsilon^2$ in the Euclidean plane), then the average sum of probabilities in an ε-ball is around C_1. Hence, an "average" station still receives the number of messages similar to those received with the original probabilities p_v. If, however, the stations are not uniformly distributed, it is still the case that stations in the smallest ball with sum of probabilities at least C_1 could receive many message after probabilities p_v are scaled up to $c_\varepsilon p_v$. On the other hand, the situation in a (very) sparse ε-ball B is as follows (by "sparse ε-ball" we mean a ball with sum of probabilities much smaller than the average):

[4] Consider for example stations $v_1, \ldots, v_n$ on a line, where the distance between v_i and v_{i+1} is $1/2^i$, for $1 \leq i < n$.

- The probability that $v \in B$ receives a message from other station from B is small (the sum is so small that usually no one is transmitting);
- The probability that $v \in B$ receives a message from $u \notin B$ is small as well (the sum in $B(v,1)$ is as large after scaling up by constant c_ε that usually the interference prevents any successful transmission from distance larger than $\varepsilon/2$).

Using this idea, our coloring algorithm works as follows. Procedure DensityTest verifies whether the sum of p_v's in a unit-disk around a station v is close to C_1 (i.e., whether v receives many messages when transmissions occur with probabilities p_v). If it is the case, procedure Playoff verifies if the density in close proximity of v is large (i.e., whether v still receives many messages when probabilities are scaled up). In the case of positive outcomes of both procedures, v is switched off, which decreases the sum of probabilities in the unit disk around v. If repeated sufficient number of times, Playoffs allow to preserve property 1 (Lemma 1), i.e., prevents the sums of probabilities in unit disks from going above C_1. On the other hand, as the positive result of Playoff cannot happen in (very) sparse areas while probabilities of stations (if not switched off) grow up to the constant $p_{\max}$, the property 2 (Lemma 2) is preserved at the end of an execution of StabilizeProbability.

```
1: procedure DensityTest(v)
2:     for c_0 log n rounds do transmit with prob. p_v
3:     if received at least c_1 log n messages then return True
4:     else return False
```

```
1: procedure Playoff(v)
2:     for c_2 log n rounds do transmit with prob. p_v · c_ε
3:     if received at least c_3 log n messages then return True
4:     else return False
```

Algorithm 1 StabilizeProbability(v)

```
1: p_v ← p_start                                   ▷ p_start = C_1/(2n)
2: p_max ← C_2/c_ε
3: while p_v < p_max do
4:     for c' times do
5:         if DensityTest(v) and Playoff(v) then
6:             v quits with color p_v
7:     p_v ← 2p_v
8: v quits with color 2p_max
```

3.2 DensityTest

In this section we fix the constants C_1, c_0 and c_1 and state properties of DensityTest which are satisfied for this choice of constants.

From now on assume that C_1 is any value such that the bounded density property holds with such C_1 that the effective communication property is satisfied as well. The goal is to choose c_0 and c_1 such that the probability of receiving a successful transmission is around c_1/c_0 in a unit ball with the sum close to $C_1/2$ and it is much lower than c_1/c_0 if the sum is significantly smaller than $C_1/2$. As possibility of sensing stations in distance close to 1 heavily depends on network topology and because of some technical reasons, the actual properties (provably) guaranteed by our choice of c_0 and c_1 will be a bit different from this intuitive goal.

Proposition 1. *Assume that bounded density property is satisfied with the parameter C_1 guarantying the effective communication property. Then, one can choose c_0, c_1 and c_d such that, for every node v, the following properties are satisfied:*

(1) Let $y = \chi(1/6, 1)$ be the number of balls of radius $1/6$ sufficient to cover a unit ball. If $1/2 \geq \sum_{w \in B(v,2/3)} p_w \geq C_1/(2y)$ then the routine DensityTest(v) returns True whp.

(2) There exists a constant c_d, such that if $\sum_{w \in B(v,1)} p_w < C_1 c_d$ then the routine DensityTest(v) returns False with high probability.

The above proposition shows that DensityTest gives an opportunity to distinguish areas with large sums of probabilities from those with much smaller sums of probabilities.

Let success in a round of DensityTest(v) means that v successfully receives or sends a message in that round.

(1) Let $C_1' = C_1/(2y)$. By the effective communication property, the probability of success is at least $\frac{1}{2} \cdot \sum_{w \in B(v,2/3)} p_w/2 \geq C_1'/4$ (cf. Fact 4). If one chooses large enough c_0 and c_1 such that $c_1/c_0 \leq \frac{1}{2} \cdot C_1'/4$, then the probability of success in a round is at least $2 \cdot \frac{c_1}{c_0}$. And, using a standard Chernoff bound, the result of DensityTest is True whp.

(2) For c_0 and c_1 chosen before, we adjust the constant c_d such that the second claim is satisfied. As a station cannot receive a message from distance larger than 1, the probability of success is not larger than the probability that at least one station from $B(v, 1)$ is transmitting in a round. This probability is bounded from above by $\sum_{w \in B(v,1)} p_w$. Therefore, if

$$\sum_{w \in B(v,1)} p_w < \frac{1}{2} \cdot \frac{c_1}{c_0} \tag{2}$$

then the result of DensityTest is False whp, by a Chernoff bound. As $c_1/c_0 \approx \frac{1}{2} \cdot C_1'/4 = C_1/(16y)$, the claim holds for $c_d < 1/(16y)$.

3.3 Proof of Lemma 1

In this section we prove Lemma 1, assuming that C_1, c_0 and c_1 are the constants satisfying properties stated in Proposition 1. Moreover, we determine values of c_2 and c_3, which depend on c_ε (and c_ε can be arbitrary at this stage).

Recall that an execution of lines 4–7 of StabilizeProbability is called a *phase*. As the initial probabilities are set to p_{start}, we have $\sum_{w \in B(v,1)} p_w \leq C_1/2$ for each v at the beginning of the algorithm. Therefore, it suffices to show that $\sum_w p_w \leq C_1$ at the end of a phase, provided the same inequality is satisfied at the beginning of this phase, where the sum is taken over w that are active in a given round. As the probabilities of active stations are multiplied by 2 at the end of each phase, the above condition for correctness of Lemma 1 can be deduced from the following lemma.

Lemma 3. *For every v, if before "For c' times" loop the following inequality holds*

$$\sum_{\substack{w \text{ is active} \\ w \in B(v,1)}} p_w < C_1$$

then after the loop

$$\sum_{\substack{w \text{ is active} \\ w \in B(v,1)}} p_w < \frac{C_1}{2}.$$

The remaining part of this section is devoted to the proof of Lemma 3, which in turn follows from the following property, provided c' is chosen large enough.

Lemma 4. *There exists a constant q which satisfies the following statement. For every v if $\sum_{w \in B(v,1)} p_w \geq C_1/2$ then there is a set of stations $S \subseteq B(v, 4/3)$ such that $\sum_{w \in S} p_w c_\varepsilon \geq q$ and every station from S is turned off in line 6 of the algorithm whp.*

Now, we prove Lemma 3 assuming correctness of Lemma 4 and then we give the proof of Lemma 4. By Lemma 4 the sum of probabilities in $B(v, 4/3)$ will decrease by at least q/c_ε in every iteration of the "For c' times" loop. The maximal sum of probabilities in $B(v, 4/3)$ is at most $\chi(1, \frac{4}{3})C_1$. Thus by performing

$c' = \chi(1, \frac{4}{3})C_1 c_\varepsilon/q$ iterations we have the "opportunity" to reduce all the probabilities in $B(v, 4/3)$, as long as $\sum_{w \in B(v,1)} p_w \geq C_1/2$. Thus after c' iterations of the loop we have $\sum_{w \in B(v,1)} p_w < C_1/2$.

It remains to prove Lemma 4; the proof is presented in the following part of this section. As we mainly analyze Playoff below, where stations transmit with probabilities scaled up by the factor c_ε, we use the notion of *mass of probability* of some set of stations A as $c_\varepsilon \sum_{w \in A} p_w$.

The proof of Lemma 4 requires to show that, close to each dense unit ball, a group of stations S with probability mass $\geq q$ exists, for which DensityTest and Playoff return true whp. The main effort in the proof is in ensuring that the elements of S can hear a message with constant probability p in each round of Playoff (i.e., when the probabilities are scaled up by c_ε) and with probability $\geq c_1/c_0$ in each round of DensityTest (i.e., with "standard" probabilities p_v's). (When this property is shown, one can adjust the value c_3/c_2 to p.) We show existence of such S by first proving that, in neighborhood of a dense unit ball, a ball with probability mass $\geq q$, center x and of radius r exists, which satisfies the following properties for some b (Lemma 5):

(a) the probability mass of each ball of radius r inside $B(x, br)$ is at most $z^\gamma q$;

(b) b is large enough to guarantee that the number of balls of radius r necessary to cover $B(x, br)$ is such that they can accumulate the whole probability mass C_1 of the unit ball, provided each of them has the (maximal) mass $z^\gamma q$, where $z > 2$ is some constant;

(c) $br \leq 1/6$;

(d) the sum of probabilities of stations in the ball of radius $2/3$ concentric with B is $\geq C_1/(2y)$, where $y = \chi(1/6, 1)$.

Using the above properties, the chances of receiving a message by a station during Playoff are estimated in Lemma 6. Then, c_2 and c_3 are chosen appropriately, to assure that Playoff returns true whp in S. On the other hand, (d) above guarantees that DensityTest returns true whp in S.

Before stating the following technical lemma sketched by (a)-(d) above, we estimate the value of b satisfying the condition (b). It is sufficient that $b^\gamma z^\gamma q \geq C_1 c_\varepsilon$ which means that $b \approx \frac{1}{z}(C_1 c_\varepsilon/q)^{1/\gamma}$ is suitable.

Lemma 5. *For every v, whenever $\sum_{w \in B(v,1)} p_w \geq C_1/2$, there exists x such that $B(x, r) \subseteq B(v, 4/3)$ and $r \leq (\frac{2q}{c_\varepsilon C_1})^{1/\gamma}$ such that if we denote*

$$D_0 = B(x, r)$$

$$D_1 = B(x, ar) \setminus D_0 \text{ for some } a \leq b$$

$$D_2 = B(x, br) \setminus (D_1 \cup D_0)$$

then

(1) The mass of probability in D_0 is at least q and at most $1/2$.

(2) For every $x' \in D_1 \cup D_2$ the mass of probability in $B(x', r)$ is bounded: $\sum_{w \in B(x',r)} p_w \leq z^\gamma q$.

(3) For all $w \in D_0$ we have $\sum_{u \in B(w, \frac{2}{3})} p_u \geq C_1/(2\chi(\frac{1}{6}, 1))$.

Proof. (of Lemma 5) Let v be an arbitrary vertex such that $\sum_{w \in B(v,1)} p_w \geq C_1/2$. Let B be a ball with radius $1/6$ included in $B(v, 1)$ with the largest mass of probability. Thus, $\sum_{w \in B} p_w \geq C_1/(2y)$ (cf. Proposition 1), where $y = \chi(1/6, 1)$. Observe that, if x located in a ball of radius $2/3$ concentric with B satisfies (1) and (2), then (3) is satisfied for x as well. Therefore, the idea of our proof is to start looking for x satisfying (1)–(3) in B, as defined above (see Fig. 1).

Let r_0 be a number satisfying the relationship $C_1/2 = q/r_0^\gamma$. That is the average probability mass of a ball of radius r_0 in $B(v, 1)$ is at least q. As B is chosen to have the largest probability mass among balls of radius $1/6$ included in $B(v, 1)$, the average probability mass of a ball of radius r_0 in B is at least q as well. Then, B includes a ball $B_0 = B(x_0, r_0)$ with probability mass $\geq q$ such that $r_0 = (2q/(C_1 c_\varepsilon))^{1/\gamma}$. If (2) is satisfied for x_0, (3) holds as well by the choice of B. On the other hand, if (2) is not satisfied for x_0, there is a ball of

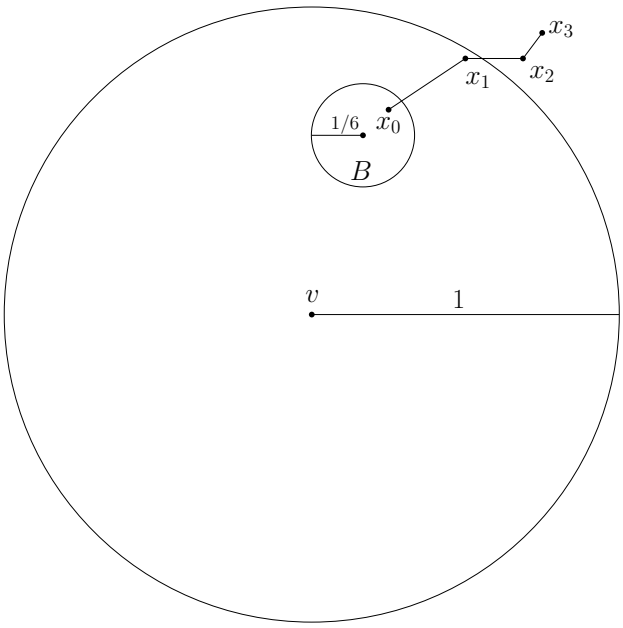

Figure 1: An illustration for the proof of Lemma 5. The distances dist(x_i, x_{i+1}) form a geometric sequence $br, br/z, br/z^2, \ldots$ B is the ball with the largest probability mass among balls of radius $1/6$ included in $B(v, 1)$.

radius r_0 and probability mass $\geq z^\gamma q$ in distance at most br_0 from x_0. Bounded growth property of the metric guarantees that this ball contains B_1 of radius $r_1 = r_0/z$ and probability mass $\geq q$.

One can build in such a way a sequence of balls $B_0, B_1, \ldots$ with probability mass larger than q, such that the radius of $B_i = B(x_i, r_i)$ is $r_i = r_{i-1}/z = r_0/z^i$ for $i > 0$, and the distance between (centers of) B_i and B_{i+1} is $\leq br_0(1 + 1/z + \cdots + 1/z^i)$, as long as (2) is not satisfied. If (2) is eventually satisfied for some x_i and $r_i = r_0/z^i$, we obtain a ball B_i of radius at most $r_0 = (2q/C_1)^{1/\gamma}$ with probability mass $\geq q$, whose center point is in distance at most $br_0 \sum_{i \geq 0} 1/z^i \leq 2br_0$ from x, provided $z > 2$. We call such an event *success*. Thus, the circle concentric with B_i of radius $2br_0 + 1/6$ contains the ball B with sum of probabilities $\geq C_1/(2\chi(\frac{1}{6}, 1))$, On the other hand,

$$br_0 \leq \frac{1}{z}(C_1 c_\varepsilon/q)^{1/\gamma} \cdot \left(\frac{2q}{c_\varepsilon C_1}\right)^{1/\gamma} \leq 1/z \leq 1/6$$

for $z \geq 6$. This implies that $2br_0 + 1/6 \leq 2/3$ and therefore (3) is satisfied as well for x_i and r_i.

It remains to show that *success* eventually appears in construction of the above sequence of balls. Note that if br_i is smaller than half of the smallest distance between stations, there is at most one station in a ball of radius br_i. This in turn implies that, if the probability mass of $B(x, r_i)$ is nonzero, then the probability mass of $B(x, br_i) \setminus B(x, r_i)$ is zero. Therefore (2) is satisfied for x and $r = r_i$. □

Lemma 6. *Let D_0 be a ball satisfying assumptions of Lemma 5 and conditions (2) and (3) stated in this lemma. Then, for every $v \in D_0$ the probability of receiving a message $p(v)$ is at least $q/8 \cdot (1/4)^{a^\gamma z^\gamma q}$.*

Proof. Now we analyze the probability that every station $v \in D_0$ receives a message. Let D denote the area around D_0, and $\mathcal{I}_0$ be the interference allowing for transmission on a distance $2r$, that is $D = D_0 \cup D_1 \cup D_2$ and by Fact 2 $\mathcal{I}_0 = \mathcal{N}/(2(2r)^\alpha)$. We also introduce four events that, when holds at the same time, allows for every station $v \in D_0$ to hear a message. Note that a station hears a message transmitted by itself.

(1) E_1 - exactly one station from D_0 transmits

(2) E_2 - no station from D_1 transmits

(3) E_3 - interference from D_2 is lower than $\mathcal{I}_0/2$

(4) E_4 - interference from stations outside D is at most $\mathcal{I}_0/2$

Observe that, since events are independent, we have

$$Pr(\text{every } v \in D_0 \text{ hears something}) \geq Pr(E_1)Pr(E_2)Pr(E_3)Pr(E_4).$$

Claim 1. $Pr(E_1) \geq q/2$

Proof. Observe that $1/2 \geq \sum_{w \in D_0} p_w \geq q$, thus by Fact 4 we have $Pr(E_1) \geq q/2$. □

Claim 2. $Pr(E_2) \geq (1/4)^{a^\gamma z^\gamma q}$

Proof. D_1 can be covered by $\chi(r, ar) \leq a^\gamma$ balls with radius r, and by Lemma 5.2 in every such ball the mass of probability is at most $z^\gamma q$, thus $\sum_{w \in D_1} p_w \leq a^\gamma z^\gamma q$. Using the inequality from Fact 5 we get $Pr(E_2) \geq (1/4)^{\sum_{w \in D_1} p_w} \geq (1/4)^{a^\gamma z^\gamma q}$. □

We choose such b that we can accommodate all the probability from an unit ball into D_2 without violating the condition (2) Lemma 5. This allows us for bounding interference in two stages. The first stage is bounding interference from close stations, and we do it more carefully as the close stations can introduce large noise. Then we bound the interference from far stations by using the fact that in every unit ball the mass of probability is at most C_1. By the properties of the metric $D_2 \cup D_1$ can be covered by b^γ balls of radius r, so we choose b as least integer that satisfy $b^\gamma z^\gamma q \geq C_1$. We also set $a = 2$ and $z = 6$ as a result of the previous observations.

Claim 3. $Pr(E_3) \geq 1/2$

Proof. Let $\mathcal{I}(D_2)$ be the interference generated by the stations from D_2. We split D_2 into *layers* $B(x, (a + i + 1)r) \setminus B(x, (a + i)r)$. Each of them can be covered by $O(i^{\gamma-1})$ balls of radius r. For the sake of clarity we assume in our calculations that the constant hidden behind the O is 1.

By bounding the interference from each layer inside D_2 we get

$$E[\mathcal{I}(D_2)] \leq \sum_{i=a-1}^{b-1} i^{\gamma-1} z^\gamma qP/(ir)^\alpha = z^\gamma qP/r^\alpha \sum_{i=a-1}^{b-1} i^{\gamma-\alpha-1} \leq z^\gamma qP/r^\alpha \sum_{i=a-1}^{\infty} i^{\gamma-\alpha-1}$$

.

Observe that, since $P = \beta\mathcal{N}$, by choosing $q \leq 1/(z^\gamma \beta 2^{\alpha+2} \sum_{i=a-1}^{\infty} i^{\gamma-\alpha-1})$ we bound the expected interference from D_2 by $\mathcal{I}_0/4$. Note that the sum $\sum_{i=a-1}^{\infty} i^{\gamma-\alpha-1}$ corresponds to the Riemann zeta function $\zeta(s)$ for $s = \alpha - \gamma + 1 > 1$, which converges to a real value. Thus, by Markov's Inequality we get $Pr(E_3) \geq 1/2$. □

Claim 4. $Pr(E_4) \geq 1/2$

Proof. We bounded the interference from stations in $B(x, 1)$ in the previous claim. Now we show that the expected outer interference is as well not too big. Let D^c denote the set of all stations outside of D, $D^c = V \setminus D$. Since in every unit ball the mass of probability is at most C_1 we can bound expected value of $\mathcal{I}(D^c)$ in the similar fashion as in the Claim 3. By splitting the space into a layers $B(x, i + 1) \setminus B(x, i)$ and covering each layer with $O(i^{\gamma-1})$ unit balls, where the probability is at most $C_1 c_\varepsilon$, we get

$$E[\mathcal{I}(D^c)] \leq \sum_{i \geq 1} i^{\gamma-1} C_1 c_\varepsilon P/(i^\alpha) \leq C_1 c_\varepsilon P \sum_{i \geq 1} i^{\gamma-\alpha-1} \leq \mathcal{I}_0/4$$

where the last inequality follows from the fact that $r \leq (\frac{2q}{c_\varepsilon C_1})^{1/\gamma}$ which implies $\mathcal{I}_0 = \mathcal{N}/(2(2r)^\alpha) \geq \mathcal{N} c_\varepsilon C_1/(2^{\alpha+4} q)$, and the choice of $q \leq 1/(2^{\alpha+4} \beta \sum_{i \geq 1} i^{\gamma-\alpha-1})$. Again, we use the Markov's Inequality to conclude the proof of the Claim. □

Finally we set $q = 1/(z^\gamma 2^{\alpha+4} \beta \sum_{i \geq 1} i^{\gamma-\alpha-1})$ with respect to the bounds from previous two claims. By combining all the claims we get $p(v) \geq Pr(\text{every } w \in D_0 \text{ hears a message}) \geq q/8 \cdot (1/4)^{a^\gamma z^\gamma q} = 2c_3/c_2$. □

Given the result of Lemma 5, we are ready to finish the proof of Lemma 4. Choose c_2 and c_3 such that $2c_3/c_2 = q/8 \cdot (1/4)^{a^\gamma z^\gamma q}$. Lemma 5 guarantees that there exists a ball D_0 satisfying conditions (1)–(3) from this lemma. Then, Lemma 6 guarantees that Playoff returns true whp, and (3) that DensityTest returns true whp for each active element of D_0. This concludes the proof of Lemma 4.

3.4 Proof of Lemma 2

Let $\varepsilon' = \varepsilon/2$. Here we prove Lemma 2 stating that for every station v the probability distributed among the stations from $B(v, \varepsilon')$ after execution of StabilizeProbability procedure is bounded from below by a constant. We show this for given C_1, c_2, c_3, we also fix the values of c_ε and C_2 to be respectively $c_\varepsilon \leftarrow 1/(\varepsilon^\alpha C_1 c_d) \cdot 8\ln(4c_2/c_3)$ and $C_2 \leftarrow \min(c_3/(8c_2), C_1 c_d/(2c_d))/c_\varepsilon$, where c_d is the constant from Prop. 1.

Intuitively, c_ε is chosen large enough to make successful transmissions on distance larger than ε' very unlikely, provided the probability mass in a considered unit ball is large enough (close to C_1). Increased probabilities generate large noise preventing communication on distance larger than ε'. Since the Playoff uses a constant probability scaled up by the factor of c_ε in the following proof we use the constant $C_2' = C_2 c_\varepsilon$ for the sake of clarity.

The main step in the proof is to show that, by our choice of C_2 and c_ε, a station whp does not quit if the mass of probability in its close proximity is small.

Lemma 7. *For every node v if*

$$\sum_{\substack{w \text{ is active} \\ w \in B(v,\varepsilon')}} p_w c_\varepsilon < 2C_2'$$

then the probability that v turns off with color p_v is negligible.

First, we show that Lemma 2 follows from Lemma 7. Then, the proof of Lemma 7 will be provided. Let one execution of lines 4-7 in StabilizeProbability be a *phase* of the algorithm. Note that all stations switched off (by quitting in line 7) during a phase have the same color (probability). Consider any v participating in an execution of the protocol. If any $w \in B(v, \varepsilon')$ does not switch off until the end of the last phase then the final value of p_w is equal to $2p_{\max} = 2C_2'/c_\varepsilon$ and therefore $\sum_{\substack{w:p_w=2p_{\max} \\ w \in B(v,\varepsilon')}} p_w c_\varepsilon \geq C_2'/c_\varepsilon$ which in turn means that the statement of Lemma 2 is satisfied. Thus, consider the case that all elements of $B(v, \varepsilon')$ switch off before the last phase. Let j be the phase in which v quits and let B_i denote the set of stations from $B(v, \varepsilon')$ active after the ith phase. Then, by Lemma 7, $\sum_{w \in B_{j-1}} p_w c_\varepsilon \geq 2C_2'$ whp. Let k be the last phase such that $\sum_{w \in B_k} p_w c_\varepsilon \geq 2C_2'$. Such phase k exists because of our assumption that all stations eventually switch off. As the probabilities of active stations are multiplied by 2 after each phase, the condition $2\sum_{w \in B_{k+1}} p_w c_\varepsilon < 2C_2'$ due to the choice of k. Thus, the set $B_k \setminus B_{k+1}$ of stations quitting in the $(k+1)$st phase satisfies $\sum_{w \in B_k \setminus B_{k+1}} p'_w c_\varepsilon \geq C_2'$, where $p'_w = 2p_w$ is the probability assigned to w in phase $k+1$. This shows that the color assigned to stations in phase $k+1$ satisfies the inequality from Lemma 2.

Now, it remains to prove Lemma 7. Observe that v turns off only when it receives at least $c_3 \log n$ messages during the PlayOff and DensityTest(v) returns True. By Proposition 1 we only need to consider the case when

(a) $\sum_{w \in B(v,\varepsilon')} p_w c_\varepsilon < 2C_2'$ and

(b) $\sum_{w \in B(v,1)} p_w \geq C_1 c_d$.

In the remaining part of the proof we show that the probability of receiving $c_3 \log n$ messages during PlayOff is negligible for our choice of C_2' and c_ε, provided (a) and (b) hold.

Let $p(v)$ denote the probability that v receives a message if all active stations transmit with currently assigned probabilities. Below, we express a condition regarding $p(v)$ which is sufficient for correctness of Lemma 2.

Fact 8. *If $p(v) < c_3/(2c_2)$ then PlayOff(v) returns False with high probability.*

Proof. The expected number of rounds in which v receives a message is $c_2 \log n \cdot p(v) \leq \frac{c_3 \log n}{2}$. Thus, by Chernoff Bound we can make the probability that PlayOff(v)=true (v receives at least $c_3 \log n$ messages) arbitrarily small by increasing c_2 and c_3 without changing the initial ratio c_3/c_2. □

In the following, let S_v be the sum of all signals received at node v, i.e.,

$$S_v = \sum_{w \text{ is transmitting}} \beta \mathcal{N}/(\text{dist}(v,w))^\alpha.$$

The next two facts show the way to bound the probability $p(v)$ of receiving a message by v in terms of the sum of probabilities in close neighborhood of v and of the interference from the whole network. These facts combined with Fact 8 give the property claimed in Lemma 2.

Fact 9. *If $S_v > \frac{2\beta\mathcal{N}}{\varepsilon'^\alpha}$ then v cannot receive a message from the outside of $B(v,\varepsilon')$.*

Proof. The strength of signal from node in a distance at least ε' from v is at most $\beta\mathcal{N}/\varepsilon'^\alpha$ (since $P=\beta\mathcal{N}$), and the interference $\mathcal{I}$ is at least $S_v - \frac{\beta\mathcal{N}}{\varepsilon'^\alpha} > \frac{\beta\mathcal{N}}{\varepsilon'^\alpha}$. From these observations we have $SINR \leq \beta\mathcal{N}/(\varepsilon'^\alpha(\mathcal{N}+\mathcal{I})) < \beta/(\varepsilon'^\alpha+\beta) < \beta$, since $\beta \geq 1$. □

Fact 10. *The probability of receiving a message at v can be bounded as follows*

$$p(v) < Pr\left(S_v \leq \frac{2\mathcal{N}\beta}{\varepsilon^\alpha}\right) + \sum_{w\in B(v,\varepsilon')} p_w \cdot c_\varepsilon$$

Proof. The former summand corresponds to the event of receiving a message from a station in a distance at least ε, thus by the Fact 9, the sum of signals at v should be at most $\frac{2\mathcal{N}\beta}{\varepsilon^\alpha}$. The latter is trivial upper bound on the probability of receiving a message from some station in $B(v,\varepsilon')$. □

Note that in the setting considered in the proof of Lemma 7 the quantity $\sum_{w\in B(v,\varepsilon')} p_w\cdot c_\varepsilon$ can be bounded by $2C_2'$. By our choice of $C_2' \leq c_3/(8c_2)$, in order to prove Lemma 7, it suffices to show that $Pr(S_v \leq \frac{2\mathcal{N}\beta}{\varepsilon^\alpha}) < c_3/(4c_2)$. Then, by Fact 10, if $\sum_{w\in B(v,\varepsilon')} p_w c_\varepsilon < 2C_2 = c_3/(4c_2)$ then $p(v) < c_3/(2c_2)$ which in turn gives the statement of Lemma 7 (by Fact 8).

From now on we focus on bounding the probability that interference at v allows for successful transmission on the distance ε or greater, i.e. that $S_v \leq \frac{2\mathcal{N}\beta}{\varepsilon^\alpha}$. As we already pointed out, it is sufficient to bound this probability from above by $c_3/(4c_2)$.

Proposition 2. *Assuming that DensityTest(v) is satisfied the following inequality holds with high probability: $Pr(S_v \leq \frac{2\mathcal{N}\beta}{\varepsilon^\alpha}) < c_3/(4c_2)$.*

Proof. In the proof, we take advantage of the fact that S_v is at least the sum of signals arriving from $B(v,1)\setminus B(v,\varepsilon')$ and, according to (a), (b) and the fact that $C_2' \leq C_1c_d/2$, the sum of probabilities in this area is at least $C_1c_d/2$. As each transmitter in Playoff uses its probability scaled up by the factor c_ε,

$$\begin{aligned} E(S_v) &\geq c_\varepsilon \textstyle\sum_{w\in B(v,1)\setminus B(v,\varepsilon')} p_w\cdot \beta\mathcal{N}/(\mathrm{dist}(w,v)^\alpha) \\ &\geq c_\varepsilon \textstyle\sum_{w\in B(v,1)\setminus B(v,\varepsilon')} p_w\cdot \beta\mathcal{N} \\ &\geq \beta\mathcal{N}c_\varepsilon C_1c_d/2 \\ &= \tfrac{2\beta\mathcal{N}}{\varepsilon^\alpha}\cdot 4\ln(4c_2/c_3) \end{aligned}$$

where the second inequality follows from the fact that $\mathrm{dist}(v,w)\leq 1$ for each $w\in B(v,1)$, the third inequality from the fact that $\sum_{w\in B(v,1)\setminus B(v,\varepsilon')} p_w \geq C_1c_d/2$.

On the other hand, the above estimation of S_v can be seen as the sum of independent random variables X_w equal either 0 or $\beta\mathcal{N}/\mathrm{dist}(v,w)^\alpha$ over all $w\in B(v,1)\setminus B(v,\varepsilon') = B$. Thus, each of these variables satisfies $X_w \leq \beta\mathcal{N}/\varepsilon^\alpha$ and $E(S_v) \geq E(\sum_{w\in B} X_w) \geq \frac{2\beta\mathcal{N}}{\varepsilon^\alpha}\cdot 4\ln(4c_2/c_3)$.

Now,we scale the variables X_w in order to apply the Chernoff bound. Let $Y_w = X_w/(\beta\mathcal{N}/\varepsilon^\alpha) \leq 1$, let $Y = \sum_{w\in B} Y_w$. Then,

$$E(Y) = E(\sum_{w\in B} X_w)/(\beta\mathcal{N}/\varepsilon^\alpha) \geq 2\cdot 4\ln(4c_2/c_3).$$

Moreover,

$$\begin{aligned} Pr(S_v \leq \tfrac{2\mathcal{N}\beta}{\varepsilon^\alpha}) &< Pr(\textstyle\sum_{w\in B} X_w \leq \tfrac{2\mathcal{N}\beta}{\varepsilon^\alpha}) \\ &= Pr(Y\leq 2) \\ &\leq Pr(Y\leq \tfrac{1}{2}E(Y)) \\ &\leq \exp(-E(Y)/8) \leq c_3/(4c_2) \end{aligned}$$

where the third last inequality follows from $4\ln(\frac{4c_2}{c_3}) \geq 2$, since $c_2 \geq c_3$. This finishes the proof of Proposition 2. □

4 Broadcast

4.1 Broadcast with non-spontaneous wakeup

Algorithm NoSBroadcast For the model with non-spontaneous wake-up, we present the algorithm NoSBroadcast in which a message is disseminated over the network in time $O(D \log^2 n)$. The algorithm works in D phases. Each phase has $O(\log^2 n)$ rounds and consists of two parts. A node participates in the phase (is *active*) if it knows the source message at the beginning of the phase. The first part of a phase executes StabilizeProbability on the set of active stations. This execution takes $O(\log^2 n)$ rounds. As a result, it assigns a color p_v to each active node v. This coloring satisfies conditions from Lemma 1 and Lemma 2. In the second part, each active node transmits the message with probability $\frac{p_v}{c\varepsilon \log n}$, for some constant c, for $O(\log^2 n)$ rounds. Consider any shortest path $s = v_0, v_1, \ldots, v_k$ in the communication graph from the source s to a node v_k. Our construction guarantees that the i-th vertex v_i of the path knows the source message after the i-th phase of the algorithm whp.

Theorem 1. *The NoSBroadcast algorithm solves the broadcast problem in the non-spontaneous wakeup model in $O(D \log^2 n)$ rounds whp.*

Theorem 1 follows directly from Fact 7 and the lemma below.

Lemma 8. *There exists a constant c such that each neighbour in the communication graph of each active node v receives the source message whp in the second part of a phase.*

Let us recall that $\varepsilon' = \varepsilon/2$. In order to satisfy the claim of Lemma 8, it is sufficient that, for each active node v, an active station v' such that $\text{dist}(v, v') \leq \varepsilon'$ transmits and is heard in distance $1 - \varepsilon'$ during NoSBroadcast (see Fact 1). We show that this is actually the case in the following proposition, which concludes the proof of Lemma 8, and thus also Theorem 1.

Proposition 3. *There exists a constant c such that for any node v active in a phase there exists a node $v' \in B(v, \varepsilon')$ which transmits in the second part of the phase and v' is heard anywhere in the distance $1 - \varepsilon'$ whp.*

Proof. The sufficient condition for occurrence of the event from the proposition in a given round is that the following three assertions hold:

(1) exactly one station transmits in $B(v, \varepsilon')$,

(2) no other station in $B(v, 2)$ transmits,

(3) the interference from outside of $B(v, 2)$ in any point of $B(v, 1)$ is smaller, than $\mathcal{I} = \mathcal{N}\alpha\varepsilon'$ (which allows hearing transmissions from the distance $1 - \varepsilon'$).

Now, in a given round, we bound from below the probabilities of the events (1)–(3) by choosing sufficiently large c.

(1) The probability, that exactly one node in $B(v, \varepsilon')$ transmits is bigger than $\frac{C_2}{2c\varepsilon \log n}$ whp. This follows from Fact 4.

(2) The probability, that no one transmits in $B(v, 2) \setminus B(v, \varepsilon')$ is bigger than $3/4$ whp.

When restricting to stations w of the same color, the inequality $\sum_{w \in B} p_w \leq C_1$ holds for any unit ball B, by Lemma 1. Thus, using the bounded growth property, we know that for each color $\sum_{w \in B(v,2)} p_w = O(1)$. As there are at most $\log n$ colors, this sum over all colors fulfills

$$\sum_{w \in B(v,2)} \frac{p_w}{c\varepsilon \log n} = \log n \frac{O(1)}{c\varepsilon \log n} = O(1).$$

Thus, if c is sufficiently large, the average number of transmitters is smaller than $1/4$, and by Markov bound no one transmits with probability $3/4$. This holds whp.

(3) The probability, that in some point of $B(v,1)$ the interference exceeds $\mathcal{I} = \mathcal{N}\alpha\varepsilon'$ is smaller than $1/4$ whp.

Once again, it is sufficient to choose c large enough so that the expected maximum of interference in $B(v,1)$ from outside of $B(v,2)$ is smaller than $\mathcal{I}/4 = \mathcal{N}\alpha\varepsilon'/4$. This follows from interference estimations similar as in Fact 6. Then this maximum interference is at most $\mathcal{I}$ with probability $3/4$ by the Markov bound.

As the events (1)–(3) are independent, the probability that exactly one node in $B(v,\varepsilon')$ transmits in a given round and it is heard in range $1-\varepsilon'$ is bigger than

$$C_2/(2c\varepsilon \log n) \cdot \frac{3}{4} \cdot \left(1 - \frac{1}{4}\right) > C_2/(4c\varepsilon \log n) \; ;$$

For further references, we state it as a separate fact.

Fact 11. *There exists a constant c such that if each node v is transmitting with probability $\frac{p_v}{c\varepsilon \log n}$ in a round, the probability that exactly one node in $B(v,\varepsilon')$ transmits and it is heard in range $1-\varepsilon'$ is at least $p = C_2/(4c\varepsilon \log n)$.*

Now, we take into account that the second part of a phase lasts for many subsequent rounds. Let the number of rounds be $T = (a \ln n)/p = O(\log^2 n)$, where a is an arbitrary constant. The probability that not all stations in range $1-\varepsilon$ from a given active node v get the message during the second part is at most

$$(1-p)^{a \ln n/p} < e^{-a \ln n} = n^{-a} \; .$$

Hence, all neighbours of v in G get the source message whp during part two of the phase. □

4.2 Broadcast with spontaneous wakeup

Algorithm SBroadcast Now, we consider the model with spontaneous wake-up. We present an algorithm SBroadcast, which starts from a *single* execution of StabilizeProbability on all stations from the network for $\varepsilon'' = \varepsilon/3$. The color assignment made during StabilizeProbability might be viewed as a kind of a communication backbone.[5] When coloring is done, the source node transmits the message deterministically. Then, after receiving the source message, every node transmits it with probability $\frac{p_v}{c\varepsilon \log n}$ in each of next $O(\log^2 n)$ consecutive rounds. (The values of c and p are the same as in Fact 11.) This assures that in each round the broadcast message is propagated by one edge, say $\{v,u\}$, of the communication graph ie. $B(v,\varepsilon'')$ informs all nodes of $B(u,\varepsilon'')$ with probability p. Using standard concentration bounds for sums of independent random variables, one can assure that the source message is delivered to each station in the graph-distance at most D from the source in $O(D \log n + \log^2 n)$ rounds whp.

Theorem 2. *The SBroadcast algorithm solves the broadcast problem in the spontaneous wake-up model in $O(D \log n + \log^2 n)$ rounds whp.*

Proof. If $B(v,\varepsilon'')$ knowing the message informs any $B(u,\varepsilon'')$ for $\{v,u\} \in G$ whp which can be proved similarly to Proposition 3. Consider a shortest path from s to some node v. A sufficient condition for v to receive a message from s is D successful transmissions on this path. In each transmission a subsequent node v_i gets the message together with its neighbourhood $B(v_i,\varepsilon'')$. By Chernoff bound all these transmissions happen whp in time $t = (2D + C \log n)/p = O(D \log n + \log^2 n)$. □

[5] In Yu et al. [21] the backbone is obtained as a connected dominating set of the communication graph, by applying techniques from geometric radio networks. However, this result holds only for a restricted family of networks and the analysis requires that stations cannot receive signals from distance larger than $1-\varepsilon$, i.e., their model assumes weak devices, and it is known that the power of that model is different than considered in our work (c.f., [16])

5 Application to Other Problems

In this section we outline how we can solve some distributed network problems other than broadcast using algorithms developed in the main part of the paper. We assume that all stations share a common global clock (i.e. they all have a common counter value assigned to each round). Each problem can be considered in adhoc setting or with some preexisting assignment of colors p_v fulfilling the conditions from Lemmas 1 and 2. In the latter case the coloring is used as a backbone.

Adhoc wake-up We study the wake-up problem as considered, e.g., in [3]. Each node in the network either wakes up spontaneously or gets activated by receiving a wake-up signal from another node. All active nodes transmit the wake-up message according to a given protocol. The running time of the protocol is the number of steps counted from the first spontaneous wake-up until all nodes become activated. Wake-up times are decided by an adversary.

In adhoc setting wake-up can be done analogously to the broadcast in which each awake station assumes that it has already received the same wake-up message. To assure the synchronization between stations, each (spontaneously) awaken station begins an execution of the protocol in the first round whose number is divisible by T, where $T = O(D \log^2 n)$ is the number of rounds of an execution of the broadcast protocol. (Recall that we assume global clock.) All the stations get this message whp after time $2T = O(D \log^2 n)$ from the first spontaneous wake-up in a network.

Wake-up with established coloring Now, we consider the wakeup problem in the setting, where all station have assigned colors (probabilities) p_v satisfying Lemma 1 and Lemma 2. Our algorithm for this setting works in two *phases*. In the first phase a new coloring q_v satisfying Lemma 1 and Lemma 2 is found for stations woken up spontaneously, i.e., by the adversary. For all other stations v we set $q_v = 0$. Then the message is broadcasted using the color $p_v + q_v$ for each station v. This procedure has running time $T = O(D \log n + \log^2 n)$. To assure the synchronization between stations, a station (woken up spontaneously) begins the protocol's execution in the first round whose number is divisible by T (i.e., a station woken up spontaneously ignores this spontaneous wakeup event until the earliest round number divisible by T and starts participating in the q_v coloring then, provided it has not received a message from other stations up to this moment).

Consensus in adhoc setting We consider consensus problem defined as follows. Each station v has some message m_v. In certain moments some stations wake-up spontaneously (i.e., stations are chosen and awaken by an adversary). In the end of the protocol all stations should agree on the same, say lexicographically smallest, message. We assume, that the set of possible messages is $\{0, 1, 2, \ldots, x\}$. The running time of the protocol is the number of steps counted from the first spontaneous wake-up until all nodes know that the protocol's execution is finished.

Our protocol makes agreement on the lexicographically smallest message using the following strategy. At the beginning of the protocol, stations woken up spontaneously perform wake-up in adhoc setting. In the last execution of StabilizeProbabilities in this wakeup, they establish some coloring p_v. Then stations that have the first bit of the message equal to 0 perform wake-up with established coloring p_v (in a limited time, as given above for the wake-up with established coloring problem) as they were woken up spontaneously. This wakeup is successful if and only if the smallest message has its first bit 0, so after this wakeup is done all stations learn the one-bit prefix of the lexicographically smallest message. This procedure can be iterated $\log x$ times for consecutive bits of messages' binary representations. In the i-th iteration stations that have the i-bit prefix of the m_v equal to the smallest $(i-1)$-bit prefix already known to everybody plus bit 0 appended, initiate wakeup with established coloring. This way all stations learn the i-th bit of lexicographically smallest message. Thus the consensus problem can be solved in time $O(D \log n \log x + \log^2 n \log x)$.

Leader election in adhoc setting We consider the leader election problem as a task of choosing one station in the whole network as the leader, assuming all stations start a protocol at the same moment.

At the beginning, all stations choose IDs from the set $\{1, \ldots, n^3\}$, independently at random, which guarantees uniqueness of IDs whp. Then, the stations perform the consensus protocol, as described above, on assigned IDs. This gives a solution in $O(D \log^2 n + \log^3 n)$ rounds.

References

[1] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Convexity and its applications in wireless networks. *J. ACM*, 59(4):18, 2012.

[2] C. Avin, Z. Lotker, F. Pasquale, and Y. A. Pignolet. A note on uniform power connectivity in the SINR model. In S. Dolev, editor, *ALGOSENSORS*, volume 5804 of *Lecture Notes in Computer Science*, pages 116–127. Springer, 2009.

[3] M. Chrobak, L. Gasieniec, and D. R. Kowalski. The wake-up problem in multihop radio networks. *SIAM J. Comput.*, 36(5):1453–1471, 2007.

[4] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. In *FOCS*, pages 492–501. IEEE Computer Society, 2003.

[5] S. Daum, S. Gilbert, F. Kuhn, and C. C. Newport. Broadcast in the ad hoc SINR model. In *DISC*, pages 358–372, 2013.

[6] M. Farach-Colton, A. F. Anta, and M. A. Mosteiro. Optimal memory-aware sensor network gossiping (or how to break the broadcast lower bound). *Theor. Comput. Sci.*, 472:60–80, 2013.

[7] L. Gasieniec, D. R. Kowalski, A. Lingas, and M. Wahlen. Efficient broadcasting in known geometric radio networks with non-uniform ranges. In G. Taubenfeld, editor, *DISC*, volume 5218 of *Lecture Notes in Computer Science*, pages 274–288. Springer, 2008.

[8] M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. In P. Fatourou and G. Taubenfeld, editors, *PODC*, pages 325–334. ACM, 2013.

[9] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local broadcasting in the physical interference model. In M. Segal and A. Kesselman, editors, *DIALM-POMC*, pages 35–44. ACM, 2008.

[10] M. M. Halldórsson and P. Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (2)*, volume 6756 of *Lecture Notes in Computer Science*, pages 625–636. Springer, 2011.

[11] M. M. Halldórsson and P. Mitra. Towards tight bounds for local broadcasting. In F. Kuhn and C. C. Newport, editors, *FOMC*, page 2. ACM, 2012.

[12] M. M. Halldórsson and P. Mitra. Wireless connectivity and capacity. In Y. Rabani, editor, *SODA*, pages 516–526. SIAM, 2012.

[13] T. Jurdzinski and D. R. Kowalski. Distributed backbone structure for algorithms in the SINR model of wireless networks. In M. K. Aguilera, editor, *DISC*, volume 7611 of *Lecture Notes in Computer Science*, pages 106–120. Springer, 2012.

[14] T. Jurdzinski, D. R. Kowalski, M. Rozanski, and G. Stachowiak. Distributed randomized broadcasting in wireless networks under the SINR model. In *DISC*, pages 373–387, 2013.

[15] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. In L. Gasieniec and F. Wolter, editors, *FCT*, volume 8070 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 2013.

[16] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in wireless networks of weak devices. In *ICALP (2)*, volume 7966 of *Lecture Notes in Computer Science*, pages 632–644. Springer, 2013.

[17] T. Jurdzinski and G. Stachowiak. Probabilistic algorithms for the wake-up problem in single-hop radio networks. *Theory Comput. Syst.*, 38(3):347–367, 2005.

[18] T. Kesselheim. Dynamic packet scheduling in wireless networks. In D. Kowalski and A. Panconesi, editors, *PODC*, pages 281–290. ACM, 2012.

[19] D. R. Kowalski and A. Pelc. Broadcasting in undirected ad hoc radio networks. *Distributed Computing*, 18(1):43–57, 2005.

[20] C. Scheideler, A. W. Richa, and P. Santi. An o(log n) dominating set protocol for wireless ad-hoc networks under the physical interference model, 2008.

[21] D. Yu, Q.-S. Hua, Y. Wang, H. Tan, and F. C. M. Lau. Distributed multiple-message broadcast in wireless ad-hoc networks under the SINR model. In G. Even and M. M. Halldórsson, editors, *SIROCCO*, volume 7355 of *Lecture Notes in Computer Science*, pages 111–122. Springer, 2012.

[22] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed local broadcasting algorithms in the physical interference model. In *DCOSS*, pages 1–8. IEEE, 2011.

[23] J. Yu, N. Wang, G. Wang, and D. Yu. Connected dominating sets in wireless ad hoc and sensor networks - a comprehensive survey. *Computer Communications*, 36(2):121–134, 2013.